

# Anomaly Detection in Dynamic Networks of Varying Size

Timothy La Fond<sup>1</sup>, Jennifer Neville<sup>1</sup>, Brian Gallagher<sup>2</sup>  
<sup>1</sup>Purdue University, <sup>2</sup>Lawrence Livermore National Laboratory  
 {tlafond,neville}@purdue.edu, bgallagher@llnl.gov

## ABSTRACT

Dynamic networks, also called network streams, are an important data representation that applies to many real-world domains. Many sets of network data such as e-mail networks, social networks, or internet traffic networks are best represented by a dynamic network due to the temporal component of the data. One important application in the domain of dynamic network analysis is anomaly detection. Here the task is to identify points in time where the network exhibits behavior radically different from a typical time, either due to some event (like the failure of machines in a computer network) or a shift in the network properties. This problem is made more difficult by the fluid nature of what is considered "normal" network behavior. The volume of traffic on a network, for example, can change over the course of a month or even vary based on the time of the day without being considered unusual. Anomaly detection tests using traditional network statistics have difficulty in these scenarios due to their **Density Dependence**: as the volume of edges changes the value of the statistics changes as well making it difficult to determine if the change in signal is due to the traffic volume or due to some fundamental shift in the behavior of the network. To more accurately detect anomalies in dynamic networks, we introduce the concept of **Density-Consistent** network statistics. These statistics are designed to produce results that reflect the state of the network independent of the volume of edges. On synthetically generated graphs anomaly detectors using these statistics show a 20-400% improvement in the recall when distinguishing graphs drawn from different distributions. When applied to several real datasets Density-Consistent statistics recover multiple network events which standard statistics failed to find, and the times flagged as anomalies by Density-Consistent statistics have subgraphs with radically different structure from normal time steps.

## 1. INTRODUCTION

Network analysis is a broad field but one of the more important applications is in the detection of anomalous or critical

events. These anomalies could be a machine failure on a computer network, an example of malicious activity, or the repercussions of some event on a social network's behavior [16, 12]. In this paper, we will focus on the task of anomaly detection in a dynamic network where the structure of the network is changing over time. For example, each time step could represent one day's worth of activity on an e-mail network. The goal is then to identify any time steps where the pattern of those communications seems abnormal compared to those of other time steps.

As comparing the communication pattern of two network examples directly is complex, one simple approach is to summarize each network using a network statistic then compare the statistics. A number of anomaly detection methods rely on these statistics [15, 2, 8]. Another method is to use a network model such as ERGM [13]. However, both these methods often encounter difficulties when the properties of the network are not static.

A typical real-world network experiences many changes in the course of its natural behavior, changes which are not examples of anomalous events. The most common of these is variation in the volume of edges. In the case of an e-mail network where the edges represent messages, the total number of messages could vary based on the time or there could be random variance in the number of messages sent each day. The statistics used to measure the network properties are usually intended to capture some other effect of the network than simply the volume of edges: for example, the clustering coefficient is usually treated as a measure of the transitivity. However, the common clustering coefficient measure is *statistically inconsistent* as the density of the network changes. Even on an Erdos-Renyi network, which does not explicitly capture transitive relationships, the clustering coefficient will be greater as the density of the network increases.

When statistics vary with the number of edges in the network, it is not valid to compare different network time steps using those statistics unless the number of edges is constant in each time step. A similar effect occurs with network models that employ features or statistics which are size sensitive: [11] show that ERGM models learn different parameters given subsets of the same graph, so even if the network properties are identical observing a smaller portion of the network leads to learning a different set of parameters.

The purpose of this work is to analytically characterize statistics by their sensitivity to network density, and offer principled alternatives that are *consistent estimators*, which empirically give more accurate results on networks with varying densities.

The major contributions of this paper are:

- We prove that several commonly used network statistics are **Density Dependent** and poorly reflect the network behavior if the network size is not constant.
- We offer alternative statistics that are **Density Consistent** which measure changes to the distribution of edges regardless of the total number of observed edges.
- We demonstrate through theory and synthetic trials that anomaly detection tests using Density Consistent statistics are better at identifying when the distribution of edges in a network has changed.
- We apply anomaly detection tests using both types of statistics to real data to show that Density Consistent statistics recover more major events of the network stream while Density Dependent statistics flag many time steps due to a change in the total edge count rather than an identifiable anomaly.
- We analyze the subgraphs that changed the most in the anomalous time steps and demonstrate that Density Consistent statistics are better at finding local features which changed radically during the anomaly.

## 2. STATISTIC-BASED ANOMALY DETECTION

A statistic-based anomaly detection method is any method which makes its determination of anomalous behavior using network statistics calculated on the graph examples. The actual anomaly detection process can be characterized in the form of a hypothesis test. The network statistics calculated on examples demonstrating *normal* network behavior form the null distribution, while the statistic on the network being tested for anomalies forms the test statistic. If the test statistic is not likely to have been drawn from the null distribution, we can reject the null hypothesis (that the test network reflects normal behavior) and conclude that it is anomalous.

Let  $G_t = \{V, E_t\}$  be a multigraph that represents a dynamic network, where  $V$  is the node set and  $E_t$  is the set of edges at time  $t$ , with  $e_{ij,t}$  the number of edges between nodes  $i$  and  $j$  at time  $t$ . The edges represent the number of interactions that occurred between the nodes observed within a discrete window of time. As the number of participating nodes is relatively static compared to the number of communications, we will assume that the node set is a constant; in time steps where a node has no communications we will treat it as being part of the network but having zero edges.

Let us define some network statistic  $S_k(G_t)$  designed to measure network property  $k$  which we will use as the test statistic (e.g., clustering coefficient). Given some set of time steps  $t_{Lx} \in \{t_{L1}, \dots, t_{Lmax}\}$  such that the set of graphs in those

times  $\{G_{t_{Lx}}\}$  are all examples of *normal* network behavior,  $S_k$  is calculated on each of these learning set examples to estimate an empirical *null* distribution. If  $t_{test}$  is the time step we are testing for abnormality, then the value  $S_k(G_{t_{test}})$  is the test statistic. Given a specified  $p$ -value (referred to as  $\alpha$ ), we can find threshold(s) that reject  $p$  percentage of the null distribution, then draw our conclusion about whether to reject the null hypothesis (conclude an anomaly is present) if the test statistic  $S_k(G_{t_{test}})$  falls out of those bounds. For this work we will use a two-tailed Z-test with a  $p$ -value of  $\alpha = 0.05$  with the thresholds  $\phi_{lower}$  and  $\phi_{upper}$ . Anomalous test cases where the null hypothesis is rejected correspond to *true positives*; normal cases where the null hypothesis is rejected correspond to *false positives*. Likewise anomalous cases where the null is not rejected correspond to *false negatives* and normal cases where the null is not rejected correspond to *true negatives*.

Deltacon [5] is an example of the statistic-based approach, as are many others [7, 8, 9, 14]. As these models also often incorporate network statistics, we will focus on the statistics themselves in this paper. Moreover, not all methods rely solely on statistics calculated from single networks: there are also *delta* statistics which measure the difference between two network examples. Netsimile is an example of such a network comparison statistic [1]. For a dynamic network, these delta statistics are usually calculated between the networks in two consecutive time steps. In practical use for hypothesis testing, these delta statistics function the same as their single network counterparts.

If the network properties being tested are not static with respect to the time  $t$  this natural evolution may cause  $S_k(G_t)$  to change over time regardless of anomalies which makes the null distribution invalid. It is useful in these instances to replace the statistic with a detrended version  $S_{*k}(G_t) = S_k(G_t) - f(t)$  where the function  $f(t)$  is some fit to the original statistic values. The paper by [6] describes how to do dynamic anomaly detection using a linear detrending function, but other functions can be used for the detrending. This detrending operation does not change the overall properties of the statistic so for the remainder of the paper assume  $S_k(G_t)$  refers to a detrended version, if appropriate.

### 2.1 Common Network Statistics

Listed here are some of the more commonly used network statistics for the anomaly detection.

#### Graph Edit Distance:

The graph edit distance (GED) [4] is defined as the number of edits that one network needs to undergo to transform into another network and is a measure of how different the network topologies are. In the context of a dynamic network the GED is applied as a delta measure to report how quickly the network is changing.

$$GED(G_t) = |V_t| + |V_{t-1}| - 2 * |V_t \cap V_{t-1}| + |E_t| + |E_{t-1}| - 2 * |E_t \cap E_{t-1}| \quad (1)$$

#### Degree Distribution Difference:

Define  $D_{i,t} = \sum_{j \neq i} e_{ij,t}$  to be the degree of  $i$ , the total number of messages to or from the node. The degree distribution is then a histogram of the number of nodes with a particular degree. Typically real-world network exhibit a power-law

degree distribution [3] but others are possible. To compare the degree distributions of  $t$  and  $t - 1$ , one option is to take the squared difference between the degree counts for each possible degree. We will call this the degree distribution difference (DD):

$$DD(G_t) = \sum_{k=1}^{\max_i(D_{i,t}, D_{i,t-1})} (\sum_i \mathbb{I}(D_{i,t} = k) - \mathbb{I}(D_{i,t-1} = k))^2 \quad (2)$$

Other measures can be used to compare the two distributions but the statistical properties of the squared distance described later extend to other distance measures.

### Weighted Clustering Coefficient:

Clustering coefficient is intended as a measure of the transitivity of the network. It is a critical property of many social networks as the odds of interacting with friends of your friends often lead to these triangular relationships. As the standard clustering coefficient is not designed for weighted graphs we will be analyzing a weighted clustering coefficient, specifically the Barrat weighted clustering coefficient (CB)[10]:

$$CB(G_t) = \sum_i \frac{1}{N * (k_i - 1) * s_i} \sum_{j,k} \frac{e_{i,j} + e_{i,k}}{2} a_{i,j} * a_{i,k} * a_{j,k} \quad (3)$$

Where  $a_{i,j} = I[e_{i,j} > 0]$ ,  $s_i = \sum_j e_{i,j}$ , and  $k_i = \sum_j a_{i,j}$ . Other weighted clustering coefficients exists but they behave similarly to the Barrat coefficient.

## 3. DENSITY DEPENDENCE

To illustrate why dependency of the network statistic on the edge count affects the conclusion of hypothesis tests, we will us first investigate statistics that are *density dependent*.

**Definition 3.1** A statistic  $S$  is *density dependent* if the value of  $S(G_t)$  is dependent on the density of  $G_t$  (i.e.,  $|E_t|$ ).

### Theorem 3.1 False Positives for Density Dependent Statistics

Let  $L = \{G_{t_L}\}$  be a learning set of graphs and  $G_{t_{test}}$  be the test graph. If  $S_k(G_t)$  is monotonically dependent on  $|E_t|$  and  $\{|E_{t_L}|\}$  is bounded by finite  $E_{lower}$  and  $E_{upper}$ , there is some  $|E_{t_{test}}|$  that will cause the test case to be rejected regardless of whether the network is an example of an anomaly with respect to property  $k$ .

PROOF. Let  $S_k(G)$  be a network statistic that is a monotonic and divergent function with respect to the number of edges in  $G$ . Given a set of learning graphs  $\{G_{t_L}\}$  the values of  $S_k(\{G_{t_L}\})$  are bounded by  $\max(S_k(\{G_{t_L}\}))$  and  $\min(S_k(\{G_{t_L}\}))$ , so the critical points  $\phi_{lower}$  and  $\phi_{upper}$  of a hypothesis test using this learning set will be within these bounds. Since an increasing  $|E_{t_{test}}|$  implies  $S_k(G_{t_{test}})$  increases or decreases, then there exists a  $|E_{t_{test}}|$  such that  $S_k(G_{t_{test}})$  is not within  $\phi_{lower}$  and  $\phi_{upper}$  and will be rejected by the test.  $\square$

If changing the number of edges in an observed network changes the output of the statistic, then if the test network

differs sufficiently in its number of edges compared to the learning examples the null hypothesis will be rejected regardless of the other network properties. As for why it is not sufficient to simply label these times as anomalies (due to unusual edge volume) the hypothesis test is designed to test for abnormality in a specific network property. If edge count anomalies also flag anomalies on other network properties, we cannot disambiguate the case where both are anomalous or just the message volume. If just the message volume is unusual, this might simply be an example of an exceptionally busy day where the pattern of communication is roughly the same just elevated. This is a very different case from where both the volume and the distribution of edges are unusual.

A second problem occurs when the edge counts in the learning set have high variance. If the statistic is dependent on the number of edges, noise in the edge counts translates to noise in the statistic values which lowers the statistical power of the test.

### Theorem 3.2 False Negatives for Density Dependent Statistics

For any  $S_k(G_{t_{test}})$  calculated on a network that is anomalous with respect to property  $k$ , if  $S_k(G_t)$  is dependent on  $|E_t|$  there is some value of the variance of the learning network edge counts such that  $S_k(G_{t_{test}})$  is not detected as an anomaly.

PROOF. Let  $S_k(G_{t_{test}})$  be the test statistic and  $\{G_{t_L}\}$  be the set of learning graphs where the edge count of any learning graph  $|E_{t_L}|$  be drawn according to distribution  $\Phi_E$ . If  $S_k(G_t)$  is a monotonic divergent function of  $|E_t|$  then as the variance of  $\Phi_E$  increases the variance of  $S_k(G_{t_L})$  increases as well. For a given  $\alpha$ , the hypothesis test thresholds  $\phi_{lower}$  and  $\phi_{upper}$  will widen as the variance increases to incorporate  $1 - \alpha$  learning set instances. Therefore, for a given  $S_k(G_{t_{test}})$  there is some value of  $\text{var}(\Phi_E)$  such that  $\phi_{lower} < S_k(G_{t_{test}}) < \phi_{upper}$ .  $\square$

With a sufficient amount of edge count noise, the statistical power of the anomaly detector drops to zero.

These theorems have been defined using a statistic calculated on a single network, but some statistics are delta measures which are measured on two networks. In these cases, the edge counts of either or both of the networks can cause edge dependency issues.

**Lemma 3.3** If a delta statistic  $S_k(G_t, G_{t-1})$  is dependent on either  $E_{min}(G_t) = \min(|E_t|, |E_{t-1}|)$  or  $E_{\Delta}(G_t) = \text{abs}(|E_t| - |E_{t-1}|)$  then Thms 3.1-3.2 apply to the delta statistic.

PROOF. For some delta statistic  $S_k(G_t, G_{t-1})$  if the statistic is dependent on  $|E_t|$ ,  $|E_{t-1}|$ , or both, Theorems 3.1 and 3.2 apply to any edge count which influences the statistic. If  $S_k(G_t, G_{t-1})$  depends on  $E_{\Delta} = \text{abs}(|E_t| - |E_{t-1}|)$  then as either  $|E_t|$  or  $|E_{t-1}|$  change the statistic produced changes leading to the problems described in theorems 3.1 and 3.2. If  $S_k(G_t, G_{t-1})$  depends on  $E_{min} = \min(|E_t|, |E_{t-1}|)$  then if both  $|E_t|$  and  $|E_{t-1}|$  increase or decrease the statistic is affected leading to the same types of errors.  $\square$

These theorems show that dependency on edge counts can lead to both false positives and false negatives from anomaly detectors that look for unusual network properties other than edge count. In order to distinguish between the observed edge counts in each time step and the other network properties, we need a more specific data model to represent the network generation process with components for each.

## 4. DENSITY INDEPENDENCE

Now that we have established the problems with density dependence, we need to define the properties that we would prefer our network statistics to have. To do this we need a more detailed model of how the graph examples were generated.

### 4.1 Data Model

Let the number of edges in any time step  $|E_t|$  be a random variable drawn from distribution  $M_E^n(t)$  in times where there is a normal message volume and distribution  $M_E^a(t)$  in times where there is anomalous message volume. Now let the distribution of edges amongst the nodes of the graph be represented by a  $N \times N$  matrix  $P_t$  where the value of any cell  $p_{ij,t}$  is the probability of observing a message between two particular node pairs at a particular time. This is a probability distribution so the total matrix mass sums to 1. Like edge count, treat this matrix as drawn from distribution  $M_P^n(t)$  in normal times and  $M_P^a(t)$  in anomalous times where  $k$  is the network parameter that is anomalous (for example, an atypical degree distribution). Any observed network slice can be treated as having been generated by a multinomial sampling process where  $|E_t|$  edges are selected independently from  $N \times N$  with probabilities  $P_t$ . Denote the sampling procedure for a graph  $G_t$  with  $F(|E_t|, P_t)$ . In the next section, we will detail how this decomposition into the count of edges and their distribution allows us to define statistics which are not sensitive to the number of edges in the network.

### 4.2 Density Consistency and Unbiasedness

In the above data model,  $P_t$  is the distribution of edges in the network, thus any property of the network aside from the volume of edges is encapsulated by the  $P_t$  matrix. Therefore, a network statistic designed to capture some network property other than edge count should be a function of  $P_t$  alone.

Let  $S_k(P_t)$  be some test statistic designed to capture a network property  $k$  of  $P_t$ , that is independent of the density of  $G_t$ . Since  $P_t$  is not directly observable we can estimate the statistic with the empirical statistic  $\hat{S}_k(G_t)$  where  $G_t$  is used to estimate  $P_t$ , with  $\hat{p}_{ij,t} = \frac{e_{ij,t}}{|E_t|}$ .

**Definition 4.1** A statistic  $S$  is **density consistent** if  $\hat{S}(G_t)$  is a consistent estimator of  $S(P_t)$ .

If  $\hat{S}_k(G_t)$  is a consistent estimator of the true value of  $S_k(P_t)$ , then observing more edge examples should cause the estimated statistics to converge to the true value given  $P_t$ . More specifically it is asymptotically consistent with respect to the true value as the number of observed edges increases. Another way to describe this property is that  $\hat{S}_k(G_t)$  has some bias term dependent on the edge count  $|E_t|$ , but the bias converges to zero as the edge count increases:  $\lim_{|E_t| \rightarrow \infty} \hat{S}_k(G_t) - S_k(P_t) = 0$ . Density consistent statistics allow us to perform

accurate hypothesis tests as long as a sufficient number of edges are observed in the networks. To begin, we will prove that the rate of false positives does not exceed the selected p-value  $\alpha$ .

#### Theorem 4.1 False Positive Rate for Density Consistent Statistics

As  $|E_t| \rightarrow \infty$  if  $\hat{S}_k(G_t = F(|E_t|, P_t \sim M_P^n(t)))$  converges to  $S_k(P_t)$ , the probability of a false positive when testing a time with an edge count anomaly  $G_{t_{test}} = F(|E_t|, P_t \sim M_P^a(t))$  approaches  $\alpha$ .

PROOF. Let all learning set graphs  $G_{t_{Lx}} = F(|E_{t_{Lx}}| \sim M_E^n(t_{Lx}), P_{t_{Lx}} \sim M_P^n(t_{Lx}))$  be drawn from non-anomalous  $|E|$  and  $P$  distributions and the test instance  $G_{t_{test}} = F(|E_{t_{test}}| \sim M_E^a(t_{test}), P_{t_{test}} \sim M_P^n(t_{test}))$  be drawn from an anomalous  $|E|$  distribution but a non-anomalous  $P$  distribution. If  $\hat{S}_k(G_t)$  is a consistent estimator of  $S_k(P_t)$ ,  $\lim_{|E_t| \rightarrow \infty} \hat{S}_k(G_t) = S_k(P_t)$ . Then as both  $|E_{t_{Lx}}|$  and  $|E_{t_{test}}|$  increase, all learning set instances and the test set instance approach the distribution  $S_k(P_t \sim M_P^n(t))$ . As any threshold is as likely to reject a learning set instance as the test instance, the false positive rate approaches  $\alpha$ .  $\square$

As the bias converges to zero, graphs created with the same underlying properties will produce statistic values within the same distribution, making the test case come from the same distribution as the null. Even if the test case has an unusual number of edges, as long as the number of edges is not too small there will not be a false positive. Density consistency is also beneficial in the case of false negatives.

#### Theorem 4.2 False Negative Rate for Density Consistent Statistics

As  $|E_t| \rightarrow \infty$  let  $\hat{S}_k(G_t = F(|E_t|, P_t^n \sim M_P^n(t)))$  converge to  $S_k(P_t^n)$  and  $\hat{S}_k(G_t = F(|E_t|, P_t^a \sim M_P^a(t)))$  converge to  $S_k(P_t^a)$ . If  $S_k(P_t^n)$  and  $S_k(P_t^a)$  are separable then the probability of a false negative  $\rightarrow 0$  as  $|E_t|$  increases.

PROOF. Let all learning set graphs  $G_{t_{Lx}} = F(|E_{t_{Lx}}| \sim M_E^n(t_{Lx}), P_{t_{Lx}} \sim M_P^n(t_{Lx}))$  be drawn from non-anomalous  $|E|$  and  $P$  distributions and the test instance  $G_{t_{test}} = F(|E_{t_{test}}| \sim M_E^a(t_{test}), P_{t_{test}} \sim M_P^a(t_{test}))$  be drawn from a non-anomalous  $|E|$  distribution but a  $P$  distribution that is anomalous on the network property being tested. Let  $\hat{S}_k(G_t)$  be a consistent estimator of  $S_k(P_t)$ . If  $S_k(P_t \sim M_P^n(t))$  and  $S_k(P_t \sim M_P^a(t))$  are separable, then

$$\lim_{|E_{t_{test}}| \rightarrow \infty} \hat{S}_k(G_{t_{test}})$$

and

$$\lim_{|E_{t_{Lx}}| \rightarrow \infty} \hat{S}_k(G_{t_{Lx}})$$

converge to two non-overlapping distributions and the probability of rejection approaches 1 for any  $\alpha$ .  $\square$

The statistical power of a density consistent statistic depends only on whether the  $P$  matrices of normal and anomalous graphs are separable using the true statistic value: as long as  $|E_t|$  is sufficiently large the bias is small enough that it is not a factor in the rate of false negatives.

A special case of density consistency is *density consistent and unbiased*, which refers to statistics where in addition to consistency the statistic is also an unbiased statistic of the true  $S_k(P_t)$ .

**Definition 4.2** A statistic  $S$  is **density unbiased** if  $\hat{S}(G_t)$  is a unbiased estimator of  $S(P_t)$ .

Unbiasedness is a desirable property because a density consistent statistic without it may produce errors due to bias when the number of observed edges is low.

### 4.3 Proposed Density-Consistent Statistics

We will now define a set of Density-Consistent statistics designed to measure network properties similar to the previously described dependent statistics, but without the sensitivity to total network edge count.

#### Probability Mass Shift:

The probability mass shift (MS) is a parallel to GED as a measure of how much change has occurred between the two networks examined. Mass Shift, however, attempts to measure the change in the underlying  $P$  distributions and avoids being directly dependent on the edge counts. The probability mass shift between time steps  $t$  and  $t-1$  is

$$MS(P_t) = \sum_{ij} (p_{ij,t} - p_{ij,t-1})^2 \quad (4)$$

The MS can be thought of as the total edge probability that changes between the edge distributions in  $t$  and  $t-1$ .

#### Probabilistic Degree Shift:

We will now propose a counterpart to the degree distribution which is density consistent. Define the probabilistic degree of a node to be  $PD(v_i) = \sum_{j \in V, i \neq j} p_{ij,t}$ . Then, let the probabilistic degree shift of a particular node in  $G_t$  be defined as the squared difference of the probabilistic degree of that node in times  $t$  and  $t-1$ . The total probabilistic degree shift (DS) of  $G_t$  is then:

$$DS(P_t) = \sum_i (\sum_{j \neq i} p_{ij,t} - \sum_{j \neq i} p_{ij,t-1})^2 \quad (5)$$

This is a measure of how much the total probability mass of nodes in the graph change across a single time step. If the shape of the degree distribution is changing, the probabilistic degree of nodes will be changing as well.

#### Triangle Probability:

As the name suggests, the triangle probability (TP) statistics is an approach to capturing the transitivity of the network and an alternative to traditional clustering coefficient measures. Define the triangle probability as:

$$TP(P_t) = \sum_{i,j,k \in V, i \neq j \neq k} p_{ij,t} * p_{ik,t} * p_{jk,t} \quad (6)$$

## 5. PROPERTIES OF NETWORK STATISTICS

Now that we have described the different categories of network statistics and their relationship to the network density we will characterize several common network statistics as density dependent or consistent, comparing them to our proposed alternatives. Table 1 summarizes our findings.

	Dependent	Consistent	Unbiased
Graph edit distance	✓		
Degree distribution	✓		
Barrat clustering		✓	
Mass shift		✓	
Degree shift		✓	
Triangle probability		✓	✓

**Figure 1: Statistical properties of previous network statistics and our proposed alternatives.**

### 5.1 Graph Edit Distance

**Claim 5.1** *GED is a density dependent statistic.*

When the edge counts of the two time steps are the same, the GED (Eq. 1) can be thought of as the difference in the distribution of edges in the network. However, the GED is sensitive to  $|E_t|$  in two ways: the change in the number of edges from  $t-1$  to  $t$ :  $E_\Delta = \text{abs}(|E_t| - |E_{t-1}|)$ , and the minimum number of edges in each time step  $E_{\min} = \min(|E_{t-1}|, |E_t|)$ . In both cases the statistic is density dependent, and in fact it *diverges* as the number of edges increases. The first case is discussed in Theorem 5.1, the second in Theorem 5.2.

**Theorem 5.1** *As  $E_\Delta \rightarrow \infty$ ,  $GED(G_t) \rightarrow \infty$ , regardless of  $P_t$  and  $P_{t-1}$ .*

PROOF. Let  $E_\Delta$  be  $\text{abs}(|E_t| - |E_{t-1}|)$ . Since the GED corresponds to  $|E_t| + |E_{t-1}| - 2 * |E_t \cap E_{t-1}|$ , the minimum edit distance between  $G_t$  and  $G_{t-1}$  occurs when their edge sets overlap maximally and is equal to  $E_\Delta$ . Therefore, as  $E_\Delta$  increases even the minimum (i.e., best case)  $GED(G_t, G_{t-1})$  also increases.  $\square$

**Theorem 5.2** *As  $E_{\min} \rightarrow \infty$ ,  $GED(G_t) \rightarrow \infty$  if  $P_t \neq P_{t-1}$ .*

PROOF. Let  $G_t = F(|E_t|, P_t^a)$  and  $G_{t-1} = F(|E_{t-1}|, P_{t-1}^b)$ , with  $P_t^a \neq P_{t-1}^b$ . Select two nodes  $i, j$  such that  $p_{ij,t}^a \neq p_{ij,t-1}^b$ . The edit distance contributed by those two nodes is  $\text{abs}(e_{ij,t} - e_{ij,t-1})$ . Let  $E_{\min}$  increase but  $E_\Delta$  remain constant. As  $E_{\min}$  increases the edit distance of the two nodes converges to  $\text{abs}(E_{\min} * p_{ij,t}^a - E_{\min} * p_{ij,t-1}^b) = E_{\min} * \text{abs}(p_{ij,t}^a - p_{ij,t-1}^b)$ . Since every pair of nodes with differing edge probabilities in the two time steps will have increasing edit distance as  $E_{\min}$  increases, the global edit distance will also increase.  $\square$

Since the GED measure is the literal count of edges and nodes that differ in each graph, the statistic is dependent on the difference in size between the two graphs. Even if the graphs are the same size, comparing two large graphs is likely to produce more differences than two very small graphs due to random chance. In addition, even when small non-anomalous differences occur between the probability distributions of edges in two time steps, variation in the edge count can result in large differences in GED.

### 5.2 Degree Distribution Difference

**Claim 5.2** *DD is a density dependent statistic.*

The degree distribution is naturally very dependent on the total degree of a network: the average degree of nodes is larger in networks with many edges. The DD measure (Eq. 2) is again sensitive to  $|E_t|$  via  $E_\Delta$  and  $E_{min}$ . In both cases the statistic is density dependent. The first case, in Theorem 5.3, shows that as  $E_\Delta$  increases, the DD measure will also increase even if the graphs were generated with the same  $P$  probabilities. The second case, in Theorem 5.4, shows that as  $E_{min}$  increases, small variations in  $P$  will increase the DD measure.

**Theorem 5.3** *As  $E_\Delta \uparrow$ ,  $DD(G_t) \uparrow$ , regardless of  $P_t$  and  $P_{t-1}$ .*

PROOF. Pick any node  $i$  in  $G_t$  and  $j$  in  $G_{t-1}$ . If  $|E_t|$  increases, and  $|E_{t-1}|$  stays the same, the expected degree of  $i$  increases, while  $j$  stays the same; likewise the inverse is true if  $|E_{t-1}|$  increases and  $|E_t|$  stays the same. Thus, as  $E_\Delta$  increases the probability of any two nodes having the same degree approaches zero, so the degree distribution difference of the two networks increases with greater  $E_\Delta$ .  $\square$

**Theorem 5.4** *As  $E_{min} \uparrow$ ,  $DD(G_t) \uparrow$  if  $P_t \neq P_{t-1}$ .*

PROOF. Let  $PD(i, G_t) = \sum_{k \neq i} p_{ik,t}$  be the probabilistic degree of node  $i$  for  $G_t = F(|E_t|, P_t)$ . Pick any node  $i$  in  $G_t$  and  $j$  in  $G_{t-1}$  such that  $PD(i, G_t) \neq PD(j, G_{t-1})$ . For a constant  $E_\Delta$ , as  $E_{min}$  increases, the expected degrees of  $i$  and  $j$  converge to  $D_{i,t} = E_{min} * \sum_{k \neq i} p_{ik,t}$  and  $D_{j,t-1} = E_{min} * \sum_{k \neq j} p_{jk,t-1}$  respectively. This means that the probability of the two nodes having the same degree approaches zero. Since every pair of nodes with differing edge probabilities in the two time steps will have unique degrees, the degree distribution difference will also increase.  $\square$

If a node has very similar edge probabilities in matrices  $P_t$  and  $P_{t-1}$ , when few edges are sampled it is likely to have the same degree in both time steps, and thus the impact on the degree distribution difference will be low. However, as the number of edges increases, even small non-anomalous difference in the  $P$  matrices will become more apparent (i.e., the node is likely to be placed in different bins in the degree distribution difference calculation), and the impact on the measure will be larger.

### 5.3 Weighted Clustering Coefficient

**Claim 5.3** *CB is a density consistent statistic.*

As shown in Theorem 5.5 below, the weighted Barrat clustering coefficient (CB, Eq. 3) is in fact density consistent. However, we will show later that the triangle probability statistic is also density unbiased, which gives more robust results, even on very sparse networks.

**Theorem 5.5** *CB( $G_t$ ) is a consistent estimator of CB( $P_t$ ), with a bias that converges to 0.*

PROOF. For  $G_t = F(|E_t|, P_t)$  the number of edges observed any pair of nodes can be represented using a multinomial distribution  $\frac{|E_t|!}{e_{ij,t}! \dots e_{yz,t}!} p_{ij,t}^{e_{ij,t}} \dots p_{yz,t}^{e_{yz,t}}$ . As  $|E_t| \rightarrow \infty$ , the rate of sampling a particular node pair  $i, j$  is  $p_{ij,t}$ , so:

$$\lim_{|E_t| \rightarrow \infty} e_{ij,t} = |E_t| * p_{ij,t}$$

Let  $CB(G_t, i)$  refer to the clustering coefficient of node  $i$ . Then, in the limit of  $|E_t|$  it converges to:

$$\begin{aligned} \lim_{|E_t| \rightarrow \infty} CB(G_t, i) &= \frac{\sum_{j,k} |E_t| p_{ij,t} + |E_t| p_{ik,t}}{2 * k_{i,lim} * \sum_{j \neq i} |E_t| p_{ij,t}} \\ &\quad * \left( \lim_{|E_t| \rightarrow \infty} \mathcal{I}[e_{ij,t} > 0] * \mathcal{I}[e_{ik,t} > 0] * \mathcal{I}[e_{jk,t} > 0] \right) \\ &= \frac{\sum_{j,k} p_{ij,t} + p_{ik,t}}{2 * k_{i,lim} * \sum_{j \neq i} p_{ij,t}} \\ &\quad * \mathcal{I}[p_{ij,t} > 0] * \mathcal{I}[p_{ik,t} > 0] * \mathcal{I}[p_{jk,t} > 0] \end{aligned}$$

where  $k_{i,lim} = \sum_{j \neq i} \mathcal{I}[p_{ij,t} > 0]$ . Since this limit can be expressed in terms of  $P_t$  alone, and CB is a sum of the clustering over all nodes, the Barrat clustering coefficient is a density consistent statistic.  $\square$

If we calculate the expectation of the Barrat clustering, we obtain:

$$\begin{aligned} E[CB(G_t, i)] &= E\left[\frac{1}{(k_i - 1) * s_i} \sum_{j,k} \frac{e_{ij,t} + e_{ik,t}}{2} a_{i,j} * a_{i,k} * a_{j,k}\right] \\ &= E\left[\frac{1}{(k_i - 1) * s_i}\right] \sum_{j,k} \frac{E[e_{ij,t}] + E[e_{ik,t}]}{2} \\ &\quad * E[a_{i,j}] * E[a_{i,k}] * E[a_{j,k}] \\ &= E\left[\frac{1}{(k_i - 1) * s_i}\right] \sum_{j,k} \frac{p_{ij,t} + p_{ik,t}}{2} \\ &\quad * E[a_{i,j}] * E[a_{i,k}] * E[a_{j,k}] \\ &= E\left[\frac{1}{(k_i - 1) * s_i}\right] \sum_{j,k} \frac{p_{ij,t} + p_{ik,t}}{2} \\ &\quad * (1 - p_{ij,t}^{|E_t|}) * (1 - p_{ik,t}^{|E_t|}) * (1 - p_{jk,t}^{|E_t|}) \end{aligned}$$

As this does not simplify to the limit of CB, it is not an unbiased estimator, and is thus density consistent but not unbiased. Other weighted clustering coefficients are also available but they have the same properties as the Barrat statistic.

### 5.4 Probability Mass Shift

**Claim 5.4** *MS is a density consistent statistic.*

As the true  $P$  is unobserved, we cannot calculate the Mass Shift (MS, Eq. 4) statistic exactly and must use the empirical Probability Mass Shift:  $\widehat{MS}(G_t) = \sum_{ij} (\hat{p}_{ij,t} - \hat{p}_{ij,t-1})^2$ . As shown in Theorem 5.6 below, the bias of this estimator approaches 0 as  $|E_t|$  and  $|E_{t-1}|$  increase, making this a density consistent statistic.

**Theorem 5.6**  *$\widehat{MS}(G_t)$  is a consistent estimator of  $MS(P_t)$ , with a bias that converges to 0.*

PROOF. The expectation of the empirical Mass Shift can be calculated with

$$\begin{aligned} E\left[\sum_{ij} (\hat{p}_{ij,t} - \hat{p}_{ij,t-1})^2\right] \\ = E\left[\sum_{ij} \left(\frac{e_{ij,t}}{|E_t|} - \frac{e_{ij,t-1}}{|E_{t-1}|}\right)^2\right] \end{aligned}$$

$$= \sum_{ij} E\left[\frac{e_{ij,t}^2}{|E_t|^2} + \frac{e_{ij,t-1}^2}{|E_{t-1}|^2} - 2 * \frac{e_{ij,t} e_{ij,t-1}}{|E_t| |E_{t-1}|}\right]$$

As the expectation of  $e_{ij,t}^2$  for any node pair  $i, j$  can be written as:

$$\begin{aligned} E[e_{ij,t}^2] &= \text{Var}(e_{ij,t}) + E[e_{ij,t}]^2 \\ &= \text{Var}(\text{Bin}(|E_t|, p_{ij,t})) + E[\text{Bin}(|E_t|, p_{ij,t})]^2 \\ &= |E_t| * p_{ij,t} * (1 - p_{ij,t}) + |E_t|^2 * p_{ij,t}^2 \end{aligned}$$

The expected empirical mass shift can then be written as:

$$\begin{aligned} \sum_{ij} E\left[\frac{e_{ij,t}^2}{|E_t|^2} + \frac{e_{ij,t-1}^2}{|E_{t-1}|^2} - 2 * \frac{e_{ij,t} e_{ij,t-1}}{|E_t| |E_{t-1}|}\right] \\ = \sum_{ij} \frac{|E_t| * p_{ij,t} * (1 - p_{ij,t})}{|E_t|^2} + \frac{|E_t|^2 * p_{ij,t}^2}{|E_t|^2} \\ + \frac{|E_{t-1}| * p_{ij,t-1} * (1 - p_{ij,t-1})}{|E_{t-1}|^2} \\ + \frac{|E_{t-1}|^2 * p_{ij,t-1}^2}{|E_{t-1}|^2} - 2 * \frac{|E_t| * p_{ij,t} |E_{t-1}| * p_{ij,t-1}}{|E_t| |E_{t-1}|} \\ = \sum_{ij} p_{ij,t}^2 - 2 * p_{ij,t} * p_{ij,t-1} + p_{ij,t-1}^2 \\ + \frac{p_{ij,t}(1 - p_{ij,t})}{|E_t|} + \frac{p_{ij,t-1}(1 - p_{ij,t-1})}{|E_{t-1}|} \\ = \sum_{ij} (p_{ij,t} - p_{ij,t-1})^2 + \frac{p_{ij,t}(1 - p_{ij,t})}{|E_t|} \\ + \frac{p_{ij,t-1}(1 - p_{ij,t-1})}{|E_{t-1}|} \end{aligned}$$

As the two additional bias terms converge to 0 as  $|E_t|$  and  $|E_{t-1}|$  increase, the empirical mass shift is a consistent estimator of the true mass shift, and is density consistent.  $\square$

We can improve the rate of convergence as well by using our empirical estimates of the probabilities to subtract an estimate of the bias from the statistic. We use the following *bias-corrected* version of the empirical statistic in all experiments:

$$\begin{aligned} \widehat{MS}^*(G_t) &= \sum_{ij} (\hat{p}_{ij,t} - \hat{p}_{ij,t-1})^2 \\ &\quad - \frac{\hat{p}_{ij,t}(1 - \hat{p}_{ij,t})}{|E_t|} - \frac{\hat{p}_{ij,t-1}(1 - \hat{p}_{ij,t-1})}{|E_{t-1}|} \end{aligned} \quad (7)$$

## 5.5 Probabilistic Degree Shift

**Claim 5.5** *DS is a density consistent statistic.*

Again, since the true  $P$  is unobserved, we cannot calculate the Degree Shift (DS, Eq. 5) statistic exactly and must use the empirical Probability Degree Shift:

$\widehat{DS}(\hat{G}_t) = \sum_i (\sum_{j \neq i} \hat{p}_{ij,t} - \sum_{j \neq i} \hat{p}_{ij,t-1})^2$ . As shown in Theorem 5.7 below, the bias of this estimator approaches 0, making this a density consistent statistic.

**Theorem 5.7**  $\widehat{DS}(G_t)$  is a consistent estimator of  $DS(P_t)$ , with a bias that converges to 0.

**PROOF.** The expectation of the empirical degree shift can be calculated with

$$\begin{aligned} E[\widehat{DS}(G_t)] &= E\left[\sum_i \left(\sum_{j \neq i} \hat{p}_{ij,t} - \sum_{j \neq i} \hat{p}_{ij,t-1}\right)^2\right] \\ &= E\left[\sum_i \left(\sum_{j \neq i} \frac{e_{ij,t}^2}{|E_t|^2} + \sum_{j \neq i} \frac{e_{ij,t-1}^2}{|E_{t-1}|^2} \right. \right. \\ &\quad + 2 * \sum_{j, k \neq i} \frac{e_{ij,t}}{|E_t|} * \frac{e_{ik,t}}{|E_t|} + 2 * \sum_{j, k \neq i} \frac{e_{ij,t-1}}{|E_{t-1}|} * \frac{e_{ik,t-1}}{|E_{t-1}|} \\ &\quad \left. \left. - 2 * \sum_{j, k \neq i} \hat{p}_{ij,t} \hat{p}_{ik,t-1}\right)\right] \\ &= \sum_i \left(\sum_{j \neq i} p_{ij,t}^2 + \sum_{j \neq i} \frac{p_{ij,t}(1 - p_{ij,t})}{|E_t|} + \sum_{j \neq i} p_{ij,t-1}^2 \right. \\ &\quad + \sum_{j \neq i} \frac{p_{ij,t-1}(1 - p_{ij,t-1})}{|E_t|} + 2 * \sum_{j, k \neq i} p_{ij,t} * \hat{p}_{ik,t} \\ &\quad \left. + 2 * \sum_{j, k \neq i} p_{ij,t-1} * \hat{p}_{ik,t-1} - 2 * \sum_{j, k \neq i} p_{ij,t} \hat{p}_{ik,t-1}\right) \\ &= \sum_i \left(\sum_{j \neq i} (\hat{p}_{ij,t} - \sum_{j \neq i} \hat{p}_{ij,t-1})^2 + \sum_{j \neq i} \frac{p_{ij,t}(1 - p_{ij,t})}{|E_t|} \right. \\ &\quad \left. + \sum_{j \neq i} \frac{p_{ij,t-1}(1 - p_{ij,t-1})}{|E_t|}\right) \end{aligned}$$

Which is density-consistent because the two additional bias terms converge to 0 as  $|E_t|$  increases.  $\square$

Since the bias converges to 0, the statistic is density consistent. By subtracting out the empirical estimate of this bias term we can hasten the convergence. We use the following *bias-corrected* empirical degree shift in our experiments:

$$\begin{aligned} \widehat{DS}^*(\hat{G}_t) &= \sum_i \left(\sum_{j \neq i} \hat{p}_{ij,t} - \sum_{j \neq i} \hat{p}_{ij,t-1}\right)^2 \\ &\quad - \sum_{j \neq i} \frac{\hat{p}_{ij,t}(1 - \hat{p}_{ij,t})}{|E_t|} + \sum_{j \neq i} \frac{\hat{p}_{ij,t-1}(1 - \hat{p}_{ij,t-1})}{|E_t|} \end{aligned} \quad (8)$$

## 5.6 Triangle Probability

**Claim 5.6** *TP is a density consistent and density unbiased statistic.*

Again, since the true  $P$  is unobserved, we cannot calculate the Triangle Probability (TP, Eq. 6) statistic exactly and must use the empirical Triangle Probability

$$\widehat{TP}(G_t) = \sum_{i,j,k \in V, i \neq j \neq k} \frac{e_{ij,t} * e_{ik,t} * e_{jk,t}}{|E_t|^3}$$

which is an unbiased estimator of the true statistic (shown below in Theorem 5.8). This means that there is no minimum number of edges necessary to attain an unbiased estimate of the true triangle probability.

**Theorem 5.8**  $\widehat{TP}(G_t)$  is an consistent and unbiased estimator of  $TP(P_t)$ .

PROOF. The expectation of the empirical Triangle Probability can be written as

$$\begin{aligned} E[\widehat{TP}(G_t)] &= \sum_{ijk} E[\hat{p}_{ij,t} \hat{p}_{ik,t} \hat{p}_{jk,t}] \\ &= \sum_{ijk} E\left[\frac{e_{ij,t}}{|E_t|} \frac{e_{ik,t}}{|E_t|} \frac{e_{jk,t}}{|E_t|}\right] \end{aligned}$$

As the number of edges on any node pair  $i, j$  can be represented with a multinomial, the expectation of each is  $|E_t| * p_{ij,t}$ . This lets us rewrite the triangle probability as

$$\begin{aligned} &= \sum_{ijk} \frac{p_{ij,t} * |E_t|}{|E_t|} \frac{p_{ik,t} * |E_t|}{|E_t|} \frac{p_{jk,t} * |E_t|}{|E_t|} \\ &= \sum_{ijk} p_{i,j} p_{i,k} p_{j,k} \end{aligned}$$

Therefore the empirical triangle probability is an unbiased estimator of the true triangle probability and is a density consistent statistic.  $\square$

## 6. EXPERIMENTS

Now that we have established the properties of density-consistent and -dependent statistics we will show the tangible effects of these properties using both synthetic datasets as well as data from real networks. The purpose of the synthetic data experiments is to show the ability of hypothesis tests using various statistics to distinguish networks that have different distributions of edges but also a random number of observed edges. The real data experiments demonstrate the types of events that generate anomalies as well as the characteristics of the anomalies that hypothesis tests using each statistic are most likely to find.


### 6.1 Synthetic Data Experiments

To validate the ability of Density-Consistent statistics to more accurately detect anomalies than traditional statistics we evaluated their performance on sets of synthetically generated graphs. Rather than create a dynamic network we generated independent sets of graphs using differing model parameters and a random number of edges. For every combination of model parameters, statistics calculated from graphs of one set of parameters (or pairs of graphs with the same parameters in the case of delta statistics like mass shift) became the null distribution, and statistics calculated on graphs from the other set of parameters (or two graphs, one from each model parameter in the case of delta statistics) became the test examples. Treating the null distribution as normally distributed, the critical points corresponding to a p-value of 0.05 are used to accept or reject each of the test examples. The percentage of test examples that are rejected averaged over all combinations of model parameters becomes the recall for that statistic.

To generate the synthetic graphs, we first sampled a uniform random variable representing the number of edges to insert into the graph. For each statistic we did three sets of trials with 1k-2k, 3k-5k, and 7k-10k edges respectively. For

	1k-2k edges	3k-5k edges	7k-10k edges
Edit Distance	0.01 $\pm$ 0.001	0.05 $\pm$ 0.01	0.21 $\pm$ 0.02
Degree Dist.	0.15 $\pm$ 0.03	0.26 $\pm$ 0.05	0.64 $\pm$ 0.06
Clust. Coef.	0.44 $\pm$ 0.05	0.78 $\pm$ 0.05	0.78 $\pm$ 0.04
Mass Shift	0.51 $\pm$ 0.05	0.77 $\pm$ 0.05	0.88 $\pm$ 0.04
Degree Shift	0.62 $\pm$ 0.07	0.63 $\pm$ 0.08	0.67 $\pm$ 0.08
Triangle Prob.	0.77 $\pm$ 0.04	0.94 $\pm$ 0.04	0.97 $\pm$ 0.03

**Figure 3: Recall when applying each statistic to flag synthetically generated anomalies. Each cell is an average over all model parameter pairs.**



**Figure 4: Comparison of Mass Shifts (a) and Edit Distances (b) produced by synthetic dataset pairs. The variance of the Edit Distances is due to the variable edge counts of the graphs, and leads to errors when distinguishing the two distributions.**

mass shift, edit distance, triangle probability, and clustering coefficient each synthetic graph was generated according to the edge probabilities of a stochastic blockmodel using accept/reject sampling to place the selected number of edges amongst the nodes. Each set of models was designed to produce graphs varying a certain property, i.e. triangle probability and clustering coefficient were applied to models with varying transitivity while mass shift and edit distance skew the class probabilities by a certain amount.

For the degree distribution and degree shift statistics, rather than using a stochastic blockmodel we assigned degrees to each node by sampling from a power-law distribution with varying parameters then using a fast Chung-Lu graph generation process to construct the network. Unlike a standard Chung-Lu process we allow multiple edges between nodes and we continue the process until a random number of edges are inserted rather than inserting edges equal to the sum of the degrees. The recalls are calculated the same way as with the stochastic blockmodels, using pairs of graphs with the same power-law degree distribution as the null set and differing degree distributions as the test set.

Figure 6.1 shows the average recall of each statistic when applied to all models of a certain range of edges. In general the more edges that are observed the more reliable the statistics are; however, the statistics we have proposed enjoy an advantage over the traditional statistics for all ranges of network sizes, and in some cases this improvement is as large as 200% or more.



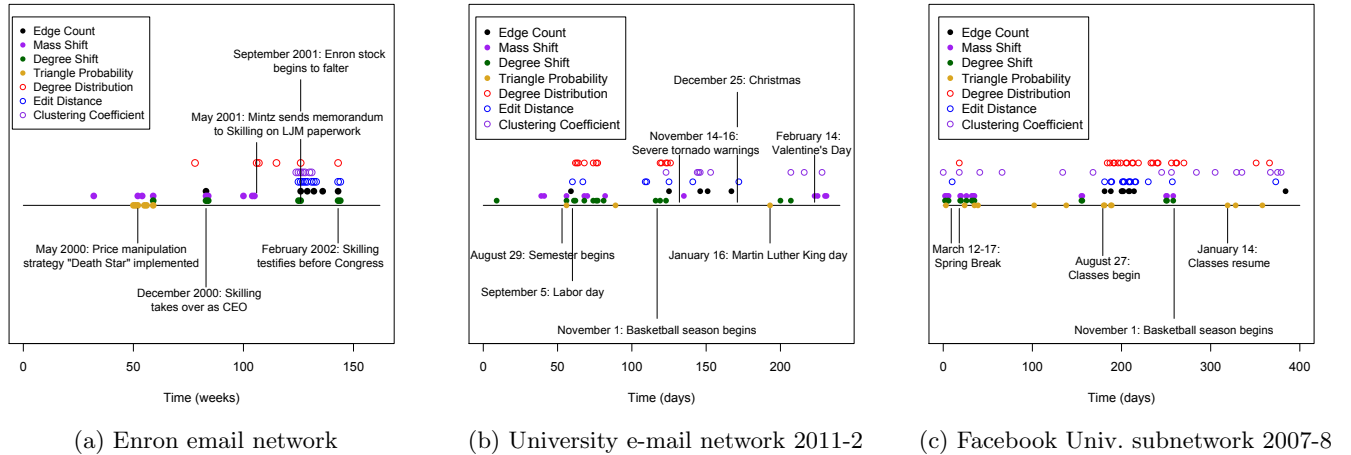


Figure 2: Timelines showing reported anomalies using each statistic for three real world networks.

Figure 4 demonstrates the effect of density consistency using one pair of graph models and the mass shift and edit distance statistics. Each black point represents a pair of datasets drawn from the same distribution while each green point is the statistic value calculated from a pair of points from differing distributions; the number of edges sampled from each distribution was 7k-10k uniformly at random. Clearly the distribution of statistic values when calculated on graph pairs from the same distribution is different from the distribution of cross-model pairs for both statistics. However, the randomness of the size of the graphs translates to variance in the edit distance statistics calculated leading to two distributions which are not easily separable leading to reduced recall. Mass shift, on the other hand, is nearly unaffected by the edge variance leading to two very distinct statistic distributions.

## 6.2 Real Data Experiments

We will now demonstrate how using Density-Consistent statistics as the test statistics for anomaly detectors improves the ability of detectors to find novel anomalies in real-world networks. We analyzed three dynamic network, one composed of e-mail communication from the Enron corporation during its operation and collapse, one from e-mail communications from university students, and one from the Facebook wall postings of the subnetwork composed of students at a university. Solid points represent statistics that we have proposed in this paper while open bubbles represent classic network statistics. Ideally the major events marked with vertical lines will be found as unusual by the detectors.

A hypothesis test using the statistic in question is applied to every time step to determine its category as anomalous or normal. The null distribution used is the set of statistics calculated on all other time steps of the stream. A normal distribution is fit to the statistic values and critical points selected set according to a p-value of 0.05; any time step with a statistic that exceeds the critical points is flagged as anomalous.


Figure 2(a) shows the the timeline of events that occurred during the Enron scandal and breakdown. The Density-Consistent

statistics are able to recover the critical events of the timeline including events that standard statistics are not able to find. In particular the price manipulation strategy that was implemented in the summer of 2000 and the CEO transition in December 2000 are not found by edge dependent statistics; unlike some of the later events this strategy was not accompanied with an abnormal number of edges so edge dependent statistics produced less of a signal on these points. In addition, the time steps that are flagged by traditional statistics cluster around a set of edge count anomalies just after the Enron stock begins to crash. This period has an elevated number of messages in general, so the statistics are responding to the number of edges in these time steps rather than changes in other properties of the network.


Figure 2(b) is taken from the e-mail communication of students from the summer of 2011 to February 2012. In general the density-consistent statistics flag time steps where certain events are taking place such as the start of basketball season, Martin Luther King Day, and Valentine's day. The density-dependent statistics tend to flag more randomly and often coincide with times of unusual total edge count.

Figure 2(c) is constructed from the set of wall postings between members of a university Facebook group. As with the Enron data, there is a large set of time steps with an usual total edge count that are also flagged by the density-dependent statistics. The density-consistent statistics recover the major events such as the beginning of each semester and the start and end of spring break. Interestingly, they also find unusual activity before and after spring break; project deadlines often fall before or after the break which may explain this activity.

Let us take a look at the network structures found in the anomalous time steps. As the mass shift, degree shift, and triangle probability statistics are sums over the edges or nodes of the graph, by decomposing the total statistic into the values generated by each edge/node we can select a subgraph which contributes the most to the statistic. As these statistics are designed to measure changes in the probability distribution of edges, this subgraph can be considered as



**Figure 5: Anomalies detected by (a) Mass Shift: Enron time steps 31-32, (b) Triangle Probability: Email time steps 88-89, and (c) Degree Shift: Facebook time steps 249-250. Each plot shows most unusual subgraph for the prior time step and the anomalous one.**



**Figure 6: Anomalies detected by (a) Edit Distance: Enron time steps 131-132, (b) Barrat Clustering: Email time steps 145-146, and (c) Degree Shift: Facebook time steps 365-366. Each plot shows most unusual subgraph for the prior time step and the anomalous one.**

the portion of the graph experiencing the most change in its edge probabilities.

Figure 5 shows the subgraphs of time steps that flag as anomalous in the mass shift, degree shift, or triangle probability; the subgraphs shown account for at least 50% of the total anomaly score. As you can see the network structures are quite different in the anomalous time steps versus the prior time steps; often nodes that do not communicate at all in the prior time step will have a significant message volume during the anomaly.

Figure 6 shows the subgraphs obtained by decomposing the mass shift, degree shift, or triangle probability on time steps that were flagged as anomalies by the density-dependent statistics but not the density-consistent statistics. Not all of the density-dependent statistics can be easily decomposed into the node/edge contributions which is why we did the density-consistent decomposition; the subgraphs generated should still be the subset of the graph experiencing the most change to its edge probabilities in that time step.

The subgraphs discovered in these time steps tend to have less dramatic changes between the anomalous and normal time step which implies that these time steps were likely flagged due to a change in the total number of edges rather than a major shift in the edge probabilities. Because density-consistent statistics are not sensitive to a global change in edge count they are better at detecting components of the network that have changed radically in their communication

behavior.

## 7. CONCLUSIONS

In this paper we demonstrate that in order to draw proper conclusions from analysis of dynamic networks it is necessary to use methods which take into account the changes they exhibit. As most network statistics were designed in an ad-hoc manner to describe informal properties of the network these statistics are not valid for use in a hypothesis testing approach to outlier detection when the network changes in size. To remedy this we have described the Density-Consistency property for network statistics and shown that statistics that adhere to this property can be used for accurate anomaly detection when the network edge volume is variable. A Density-Consistent statistic should measure some property of the network that is independent of the edge volume like the distribution of edges or the transitivity.

We have also proposed three network statistics, Mass Shift, Degree Shift, and Triangle Probability to replace the edge dependent statistics of Graph Edit Distance, Degree Distribution and Clustering Coefficient. We have proven that our statistics are Density-Consistent and demonstrated using synthetic trials that anomaly detectors using the consistent statistics have superior performance. When applied to real datasets, anomaly detectors utilizing Density-Consistent statistics are often able to recover events that are missed by edge dependent statistics.

## 8. REFERENCES

- [1] M. Berlingerio, T. D. Koutra, and C. Faloutsos. Network similarity via multiple social theories. *In Proceedings of the 5th IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 1439–1440, 2012.
- [2] K. C. D. Banks. Models for network evolution. *Journal of Mathematical Sociology*, 21:1-2:173–196, 1996.
- [3] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. *ACM SIGCOMM Computer Communication Review*, 29(4), 1999.
- [4] X. Gao, B. Xiao, D. Tao, and X. Li. A survey of graph edit distance. *Journal of Pattern Analysis and Applications*, 13(1):113–129, 2010.
- [5] D. Koutra, J. Vogelstein, and C. Faloutsos. Deltacon: A principled massive-graph similarity function. *SDM*, 2013.
- [6] T. LaFond, J. Neville, and B. Gallagher. Anomaly detection in networks with changing trends. *ODD<sup>2</sup>’14*, 2014.
- [7] I. McCulloh. Detecting changes in a dynamic social network. Technical Report CMU-ISR-09-104, Carnegie Mellon University, 2009.
- [8] I. A. McCulloh and K. M. Carley. Social network change detection. Technical Report No. CMU-ISR-08-116, Carnegie Mellon University, 2008.
- [9] e. a. Priebe, Carey E. Scan statistics on enron graphs. *Computational & Mathematical Organization Theory*, 11.3:229–247, 2005.
- [10] e. a. Saramaki, Jari. Generalizations of the clustering coefficient to weighted complex networks. *Physical Review*, 75.2, 2007.
- [11] C. R. Shalizi and A. Rinaldo. Consistency under sampling of exponential random graph models. *The Annals of Statistics*, 41:2:508–535, 2013.
- [12] V. A. Siris and F. Papagalou. Application of anomaly detection algorithms for detecting syn flooding attacks. *Computer Communications*, 29:1433–1442, 2006.
- [13] e. a. Snijders, Tom. New specifications for exponential random graph models. *Sociological methodology*, 36.1:99–153, 2006.
- [14] S. B. Snijders. Non-parametric standard errors and tests for network statistics. *CONNECTIONS*, 22(2):161–170, 1999.
- [15] T. A. Snijders. Models for longitudinal network data. *Models and methods in social network analysis 1*, pages 215–247, 2005.
- [16] A. Tartakovsky, H. Kim, B. Rozovskii, and R. Blažek. A novel approach to detection of ‘denial-of-service’ attacks via adaptive sequential and batch-sequential change-point detection methods. *IEEE Transactions on Signal Processing*, 54:9, 2006.